

We claim:

1. A method of placing content into an image sequence, comprising:
applying a transform function to visual content to be inserted into an image sequence; and
blending the transformed visual content with the image sequence;
wherein the transforming step and the blending step together result in insertion of the content into the image sequence such that the content appears at a target location as if it had been part of the original scene displayed by the image sequence; and
wherein the transforming step comprises all applied geometric transformations of the visual content, is performed without reference to any content of the image sequence, and is performed without reference to any three dimensional real world space locations or dimensions.
2. The method of claim 1 wherein the three dimensional real world space locations include a location of a camera which originally produced the image sequence.
3. The method of claim 1 wherein the three-dimensional real world scale dimensions comprise units of physical measurement relating to the scene of the image sequence.
4. The method of claim 1 wherein the blending step is performed without reference to any three-dimensional real world space locations.

5. The method of claim 1, wherein the target location may change relative to the scene as the image sequence progresses.

6. The method of claim 1, wherein the image sequence is received from an analog video recorder, a digital video recorder or a broadcast quality video camera.

7. The method of claim 1:

wherein a portion of the content is placed into an input image within the image sequence; comprising the additional steps of:

referencing a dynamic scene component model set which contains a three dimensional model for each of a subset of objects present in the scene that is displayed by the image sequence;

receiving camera sensor data corresponding to the input image;

selecting one scene component model out of the scene component model set; and

filling a graphics frame buffer (GFB) with updated pixel data; and

wherein the transforming step comprises:

transforming the selected scene component model such that, when inserted into the input image, the selected scene component model will appear as if it had been part of

the original scene displayed by the input image, thus creating an input embedded model; and

texturing the input embedded model with said GFB, thus creating an input embedded GFB model; and

wherein the blending step comprises blending the input embedded GFB model with the input image such that the input embedded GFB model appears as if it had been part of the original scene displayed by the input image, thus creating a blended image; and

wherein the method also comprises the step of transmitting the input image or the blended image according to a selection of a user.

8. The method of claim 7, wherein the sensor data relates to a camera that has collected the input image and includes:

pan information;

tilt information;

zoom information; and

focus information.

9. The method of claim 7, wherein each scene component model within the scene component model set comprises a two dimensional quadrilateral.

10. The method of claim 7, wherein:

each scene component model within the scene component model set comprises a three dimensional mesh of polygons; and

each polygon within the three dimensional mesh comprises a quadrilateral.

11. The method of claim 7, wherein:

each model in the scene component model set is represented within a normalized scale of measurement;

the scale is not required to use actual real world units of measurement;

the subset of objects includes all placement objects upon which the GFB may possibly be textured;

the subset of objects further includes all containing objects which enclose movement of any of the placement objects; and

each model in the scene component model set comprises one or more polygonal meshes.

12. The method of claim 7, wherein the transforming step further comprises:

identifying a pre-determined scene component model insert region corresponding to the target location, wherein the insert region has been pre-selected by a user;

identifying a baseline perspective transform for the selected scene component model; transforming, using the baseline perspective transform, three dimensional coordinates in scene component model space to two dimensional coordinates within a view space of a baseline image;

mapping, using the baseline perspective transform, the selected scene component model into the scene component model insert region within the baseline image, thus creating a baseline embedded model;

deriving, based solely on the sensor data and a pre-determined inductive camera model, a set of parameters for a parameterized inductive transform;

converting, using the parameterized inductive transform, two dimensional coordinates in the view space of the baseline image to two dimensional coordinates within a view space of the input image; and

transforming, using the parameterized inductive transform, the baseline embedded model such that the baseline embedded model is correct for the input image, thus creating an input embedded model.

13. The method of claim 12, wherein the scene component model insert region comprises one or more two dimensional quadrilaterals.

14. The method of claim 7, wherein the blending step comprises:

determining a set of bounding polygons which define one or more regions of interest in the input image;

performing pre-processing on pixels within the one or more regions of interest, wherein the pre-processing includes application of one or more of a set of pre-determined filters, thus creating a pre-processed input image;

performing color separation via reference to the pre-processed input image, according to a pre-determined color separation model, on the pixels within the one or more regions of interest, such that pixels that are determined to be background are given a background appropriate alpha value, while pixels that are determined to be foreground are given a foreground appropriate alpha value, thus creating a set of alpha polygons;

copying the set of alpha polygons into a previously empty image, thus creating a color separated alpha frame;

performing post-processing on the color separated alpha frame, on the pixels within the regions of interest, wherein the post-processing includes application of one or more of a set of pre-determined filters, thus creating a post-processed alpha frame; and

alpha-blending the input embedded GFB model into the input image, using a combination of the post-processed alpha frame and an alpha channel of the input embedded GFB model for said blending, thus creating the blended image.

15. The method of claim 14 wherein the alpha-blending step uses a dynamically determined blending alignment offset for two dimensional refinement of an alignment of the input embedded GFB model and the pre-processed input image.

16. The method of claim 14, wherein the color separation model has been determined by a method comprising:

- receiving a reference image sequence;
- receiving selections from a user, via an input means, of one or more regions of interest in the image sequence;
- determining, for the regions of interest, a set of color regions based on contained pixel colors;
- refining, for the regions of interest, via manual user modifications, the color region, thus creating a set of refined color regions;
- adding, for the regions of interest, each color within the set of refined color regions to a background inclusion table; and
- maintaining the background inclusion table in computer software memory such that colors are directly indexed;

wherein the background inclusion table comprises the set of colors which will be considered background within the color separation model.

17. The method of claim 7:

wherein the transmitting step comprises determining a position of an on/off switch; wherein a switch position set to on indicates transmission of the blended image; and wherein a switch position set to off indicates transmission of the input image.

18. A method of placing a virtual marker line into an image sequence, comprising:

applying a transform function to a marker line that is to be inserted into an image sequence that was produced by a camera; and

blending the transformed line with the image sequence;

wherein the transforming step and the blending step together result in insertion of the line into the image sequence such that the line appears at a target location as if it had been part of the original scene displayed by the image sequence; and

wherein the transforming step comprises all applied geometric transformations of the line, is performed without reference to any real-time content of the image sequence, and is performed without reference to any location of the camera.

19. The method of claim 18:

wherein an instance of the marker line is placed into an input image within the image sequence;

further comprising the steps of:

referencing a dynamic scene component model set which contains a three dimensional model for each of a subset of objects present in the physical scene that is displayed by the image sequence;

receiving camera sensor data corresponding to the input image;

selecting one scene component model out of said scene component model set; and

filling a graphics frame buffer (GFB) with updated pixel data; and

wherein the transforming step comprises:

transforming the selected scene component model such that, when inserted into the input image, the selected scene component model will appear as if it had been part of the original scene as displayed by the input image, thus creating an input embedded model; and

texturing the input embedded model with said GFB, thus creating an input embedded GFB model; and

wherein the blending step comprises blending the input embedded GFB model with the input image such that the input embedded GFB model appears as if it had been part of the original scene displayed by the input image, thus creating a blended image; and

wherein the method also comprises the step of transmitting the input image or the blended image according to a selection of a user.

20. The method of claim 19, wherein the scene component model set includes:

a first model of a three dimensional surface representing a first section of a playing field that is substantially bounded by two sidelines, a first goal line or end point and a midpoint inclusive; and

a second model of a three dimensional surface representing a second section of the playing field that is substantially bounded by the two sidelines, the midpoint and a second goal line or end point.

21. The method of claim 19:

wherein the selecting step comprises:

periodically receiving a user-selected value of a field location, wherein the value is not received via a direct input on an image taken from a live camera video stream; and

choosing a scene component model from the first model and the second model such that field position corresponding to the user-selected value is physically contained by the physical space represented by the selected scene component model, thus yielding the selected scene component model; and

wherein the filling step comprises:

mapping the user-selected value to a quadrilateral within a two-dimensional selection canvas, wherein the quadrilateral has a location within the selection canvas proportional to a physical location of a position on the playing field and a size that is indicative of a desired thickness of the virtual marker line, according to a user selection;

filling the quadrilateral with a desired pattern and color, according to a user selection; and

copying contents of the selection canvas into the GFB.

22. A method of placing a virtual sign into an image sequence, comprising:
 - applying a transform function to a virtual sign that is to be placed into an image sequence that was produced by a camera; and
 - blending the transformed sign with the image sequence;

wherein the transforming step and the blending step together result in insertion of the sign into the image sequence such that the sign appears at a target location as if it had been part of the original scene displayed by the image sequence; and

wherein the transforming step comprises all applied geometric transformations of the sign, is performed without reference to any content of the image sequence, and is performed without reference to location of the camera.

23. The method of claim 22:

wherein an instance of the sign is placed into an input image within the image sequence; further comprising the steps of:

referencing a dynamic scene component model set which contains a three dimensional model for each of a subset of objects present in the physical scene that is displayed by the image sequence;

receiving camera sensor data corresponding to the input image;

selecting one scene component model out of the scene component model set; and

filling a graphics frame buffer (GFB) with updated pixel data; and

wherein the transforming step comprises:

transforming the selected scene component model such that, when inserted into the input image, the selected scene component model will appear as if it had been part of

the original scene as displayed by the input image, thus creating an input embedded model; and

texturing the input embedded model with the GFB, thus creating an input embedded GFB model; and

wherein the blending step comprises blending the input embedded GFB model with the input image such the said input embedded GFB model appears as if it had been part of the original scene displayed by the input image, thus creating a blended image; and

wherein the method also comprises the step of transmitting the input image or the blended image according to a selection of a user.

24. The method of claim 23, wherein the scene component model set includes:

a first model of a three dimensional surface representing the shape of one area upon which the sign may be placed; and

a second model of a three dimensional surface representing the shape of another area upon which the sign may be placed.

25. The method of claim 23:

wherein the selecting step comprises:

periodically receiving a user-selected location for the sign, wherein the location was selected via a direct input on an image captured from the image sequence; choosing a scene component model from the first and second models depending on a physical shape of the location, thus yielding a selected scene component model; and wherein the filling step comprises copying the sign into the GFB.

26. A method of placing content on a moving object in an image sequence, comprising:
 - applying a transform function to visual content to be inserted into an image sequence, wherein the image sequence includes a moving object; and
 - blending the transformed visual content with the image sequence; wherein the transforming step and the blending step together result in insertion of the content into the image sequence such that the content appears at a target location as if it had been part of the original scene displayed by the image sequence, and such that the content is located on the moving object as the object moves in the scene; and
 - wherein the transforming step comprises all applied geometric transformations of the visual content, is performed without reference to any content of the image sequence, and is performed without reference to any three dimensional real world space locations or dimensions.
27. The method of claim 26:

wherein a portion of the content is placed into an input image within the image sequence; comprising the additional steps of:

referencing a dynamic scene component model set which contains a three dimensional model for each of a subset of objects present in the scene that is displayed by the image sequence;

receiving camera sensor data corresponding to the image;

selecting one scene component model out of the scene component model set; and

filling a graphics frame buffer (GFB) with updated pixel data; and

wherein the transforming step comprises:

transforming the selected scene component model such that, when inserted into the input image, the selected scene component will appear as if the selected scene component model had been part of the original scene displayed by the input image, thus creating an input embedded model; and

texturing the input embedded model with said GFB, thus creating an input embedded GFB model; and

wherein the blending step comprises blending the input embedded GFB model with the input image such that the input embedded GFB model appears on the moving object as if the input embedded GFB model had been part of the original scene displayed by the input image, thus creating a blended image; and

wherein the method also comprises the step of transmitting either the input image or the blended image, according to a selection of a user.

28. The method of claim 27, wherein the scene component model set includes:

a first model of a three dimensional surface representing a first area, the first area being the part of the moving object upon which the content will be placed;

a second model of a three dimensional surface representing a second area, the second area being that within which the moving object will travel; and

a third model, comprising a temporary three dimensional surface that is constructed from the first model, then oriented and positioned in real-time based on a dynamic location of the moving object within the space of the second model.

29. The method of claim 28, wherein the selecting step comprises:

periodically specifying, using an object positional tracker, a dynamic location and orientation of the moving object, wherein the dynamic location and orientation are specified within the model without reference to any location of any camera that produced the image sequence;

constructing the third model from the first model;

positioning and orienting the third model within the second model according to the dynamic location and orientation of the object; and
selecting the third model, thus yielding the selected scene component model.

30. A method of inserting dynamically changing visual content into an image sequence, comprising:

applying a transform function to dynamically changing visual content; and
blending the transformed dynamically changing visual content with an image sequence;
wherein the transforming step and the blending step together result in insertion of the content into the image sequence such that the dynamically changing visual content appears as if it had been part of the original scene displayed by the image sequence;
wherein said transforming step comprises all applied geometric transformations of the dynamically changing visual content, is performed without reference to any real-time image content of the image sequence and is performed without reference to any three dimensional real world space locations as related to the image sequence.

31. The method of claim 30:

wherein a portion of the dynamically changing visual content is placed into an input image within the image sequence; and

comprising the additional steps of:

referencing a dynamic scene component model set which contains a three dimensional model for each of a subset of objects present in the scene that is displayed by the image sequence;

receiving camera sensor data corresponding to the input image;

selecting one scene component model out of the scene component model set; and

filling a graphics frame buffer (GFB) with updated pixel data; and

wherein the transforming step comprises:

transforming the selected scene component model such that, when inserted into the input image, the selected scene will appear as if it had been part of the original scene displayed by the input image, thus creating an input embedded model; and

texturing the input embedded model with said GFB, thus creating an input embedded GFB model; and

wherein the blending step comprises blending the input embedded GFB model with the input image such that the input embedded GFB model appears as if the input embedded GFB model had been part of the original scene displayed by the input image, thus creating a blended image; and

wherein the method also comprises the step of transmitting either the input image or the blended image, according to a selection of a user.

32. The method of claim 31, wherein the scene component model set includes:

 a first model of a three dimensional surface representing the shape of a first area upon which the dynamically changing visual content may be placed; and

 a second model of a three dimensional surface representing the shape of a second area upon which the dynamically changing visual content may be placed.

33. The method of claim 32:

 wherein the selecting step comprises:

 periodically receiving a user-specified location value for the dynamically changing visual content, wherein the value has been received via a direct input on an image from the image sequence; and

 choosing a scene component model from the first model and the second model, depending on the physical shape of the user-specified location, thus yielding the selected scene component model;

 wherein the dynamically changing visual content has been rendered by a general purpose graphics rendering device; and

 wherein the filling step comprises copying the dynamically changing visual content into the GFB.

34. A method of transforming a first point P[A] within a first two-dimensional camera view to a second point P[B] within a second two-dimensional camera view, such that the point P[A] corresponds to the same location within real world space as the point P[B], comprising:

applying a transform function to transform a first point P[A] within a first image to a second point P[B] within a second image;

wherein the first image contains content related to a first camera view;

wherein the second image contains content related to a second camera view;

wherein the point P[A] corresponds to the same location within real world space as point P[B];

wherein the transforming step is performed without reference to any content of the first image and the second image;

wherein the transforming step is further performed without reference to any three-dimensional real world space locations or dimensions as related to the content of the first image and the second image.

35. The method of claim 34, wherein the transform function uses parameters comprising:

a set of fixed transform parameters;

first sensor data received from a first camera corresponding to the first camera view;

second sensor data received from a second camera corresponding to the second camera view;
a set of dynamic transform parameters that are based on the fixed transform parameters, the first sensor data and the second sensor data.

36. The method of claim 35, wherein the first and second camera sensor data include:

pan information;
tilt information;
zoom information; and
focus information.

37. The method of claim 35, wherein the fixed transform parameters have been determined by a method comprising:

receiving an image sequence from a live camera that is outfitted with camera sensors;
receiving a stream of camera sensor data from the attached camera sensors, wherein the stream of camera sensor data comprises camera sensor information applicable to images within the image sequence;

receiving selections from a user, via direct input on one or more of the images, of multiple points of interest within 2D camera view space, wherein the selections have been received at multiple levels of camera zoom and focus;

correlating, in time, each of the images for which selections were made and each of the user selections to the stream of camera sensor data;

determining, based on the user selections and the camera sensor data, a set of zoom/focus transform parameters to be included within the set of fixed transform parameters;

receiving selections from a user, via direct input on one or more images, of multiple points of interest within two-dimensional camera view space; wherein the selections have been received at multiple levels of camera pan and camera tilt;

correlating, in time, each of the images for which selections were made and each of the user selections to the first and second camera sensor data; and

determining, based on the user selections, and the camera sensor data, a set of other parameters to be included within the set of fixed transform parameters.

38. The method of claim 37, wherein the first and second camera sensor data includes:

pan information;

tilt information;

zoom information; and

focus information.